

Cornell University

AD-A267 387



DTIC  
ELECTE  
AUG 04 1993  
S A D

DEPARTMENT OF COMPUTER SCIENCE

41501 PERSON HALL  
ITHACA NY 14853-7501  
telephone (607) 255-7346  
telex W116718054  
FAX (607) 255-7428

2

April 30, 1993

Dear Madam/Sir:

Enclosed is a copy of the Third Quarterly R&D Status Report produced by the Isis Group under DARPA/ONR contract number N00014-92-J-1866. This report covers the time period between April 1993 and June 1993.

Respectfully yours,

*Maureen Robinson*

Maureen Robinson  
Administrative Aide  
(607) 255-9198

This report is unclassified and may be  
distributed without permission

This document has been approved  
for public release and sale; its  
distribution is unlimited.

93-17394



93 8 3 0 0 1

# Redesigned Isis and Meta System under Mach

## Third Quarterly R & D Status Report July 1993

Prof. Kenneth P. Birman  
Department of Computer Science  
Cornell University, Ithaca New York  
607-255-9199

Prof. Keith Marzullo  
Dept. of Computer Science  
U.C. San Diego, San Diego, California

Accession For	
NTIS	CRASH
DTIC	IAS
Univ. of	Cal.
Just	State
By	
Distribution	
Availability Codes	
Dist	Availability Codes
A-1	

DTIC QUALITY INSPECTED 3

This work was sponsored by the Defense Advanced Research Projects Agency (DoD), under contract N00014-92-J-1866 issued by the Office of Naval Research. The view, opinions and findings contained in this report are those of the authors and should not be construed as an official DoD position, policy, or decision.

## Personnel

- Academic Staff:

- Prof. Kenneth P. Birman, Principle-Investigator
- Prof. Keith Marzullo, Co-Investigator
- Dr. Robert Cooper, Research Associate
- Dr. Robbert van Renesse, Research Associate
- Dr. Aleta Ricciardi, Post-doctoral Research Associate
- Dr. Mark Wood, Post-doctoral Research Associate

- Graduate Students:

- Lorenzo Alvisi (Marzullo)
- Navin Budhiraja (Marzullo)
- Brad Glade (Birman)
- Guernsey Hunt (Birman)
- Neil Jain (Birman)
- David Karr (Birman)
- Michael Kalantar (Birman)
- Michael Reiter (Birman)
- Aleta Ricciardi (Birman)
- Laura Sabel (Marzullo)

## The Isis Project Progress

This quarter has seen solid progress toward our major project goals. Horus, the complete redesign of the original ISIS system, is now stable and in use. The design of how to effectively run Horus over high-speed ATM networks and parallel computers such as the Intel Paragon continues, and we have set up a laboratory at the University of California for the continued development of the realtime group programming toolkit Corto [ref 4].

In this quarter, we completed the first non-prototype version of the Lomita monitoring and control language [ref 9]. The lack of such a language has been a major stumbling block in the use of Meta. While the native Meta programming language is efficient, it is almost impossible to use because of its postscript-like syntax. For example, the following is a simple Lomita control rule that has a machine leave one group and join another group when told to do so (by setting a local register). This rule is taken from a Lomita program used to test Meta's group mechanism:

```
machine[*]
  whenever switch
  do LEAVE machineset[group]
    JOIN machineset[! group]
    group = !group
    switch = FALSE
    PRINT(CTIME(localtime),
          ": switched to machine group ", ! group, ".\n")
  else if STOP do
    STOP = FALSE
    PRINT("Terminated.\n");
```

Here is the same rule, in the "internal" Meta programming language, from which the user is now insulated:

```
RI2 @ GUARD "RIO RI2 @ AND localtime @ AND TGUARD
\"machineset(" RI2 STRING + ")\n" 'LEAVE \"machineset(" +
RI2 NOT STRING + ")\n" 'JOIN RI2 NOT 'WI2 0 'WIO localtime
CTIME \"\": switched to machine group \" RI2 NOT \".\n\n\"
'PRINT ALTERNATE RI7 GUARD 0 'WI7 \"Terminated.\n\n\" 'PRINT
'EXIT\" + 'LNPL 'EXITO
```

As one would expect, having the ability to now write complex Meta control rules has resulted in the identification of some bugs in Meta. We have been working through these bugs and plan to issue a new (free) university release once the system restabilizes.

We have been very active this quarter in terms of publishing. A book that collects the major ISIS papers is nearing completion. The debate in the academic community continues on the value of virtual synchrony and the underlying order-based communication protocols [refs. 2, 3, 5, 8] versus other methods of programming distributed systems. This debate, in part, has arisen from the success of ISIS in the commercial arena and the creation of various research projects working on ISIS-like systems.

## Transitions

The external user base for Isis, Horus and Meta continues to grow, including both university and industrial sites. We have recently been talking with a research group in Siemens about the use of Meta and Lomita in an industrial prototype. We have also set up an arrangement for cooperative research with a group in the Software Engineering Institute and with INESC for real-time transports and scheduling support in Corto. Through Isis Distributed Systems Inc, Isis is now installed at hundreds of locations worldwide. A good recent example of an Isis application that also uses Mach is the air traffic control system of a major European government, which has selected Isis on OSF/1 Alpha Workstations from Digital Equipment as the basis for its next generation ATC system effort, with technology deployment to begin in late 1994.

### Third Budget Statement

- a. ARPA Order Number: 9247
- b. Contract Number: N00014-92-J-1866
- c. Agent: ONR
- d. Contract Title: A Redesigned ISIS and Meta System Under Mach
- e. Organization: Cornell University
- f. PIs: Kenneth P. Birman and Keith Marzullo
- g. Actual Start Date: 9/30/92
- h. Expected End Date: 12/30/95
- i. Expected End Date if Options Exercised: NA
- j. Total Price: \$3,137,518
- k. Spending Authority Provided So Far: \$1,281,331
- l. Expenditures through 6/93 \$710,833
- m. Date When These Funds Will Be Fully Expended: 10/31/93
- n. Additional Funds Expected Per Contract (by FY):  
FY94 \$928,050  
FY95 \$928,137

## Publications

- Nonblocking and Orphan-Free Message Logging Protocols. Lorenzo Alvisi, Bruce Hoppe and Keith Marzullo. 23rd Fault-Tolerant Computing Symposium, June 1993, pp. 145-154.

Describes another approach to fault-tolerant services based on message logging protocols. The protocol here, which has been implemented, is both efficient in normal usage and in failure recovery.

- Understanding Partitions and the 'No Partition' Assumption. Andre Schiper, Aletta Ricciardi and Kenneth Birman. To appear 4th Future Trends in Distributed Computing Systems, September 22-24, 1993, Lisbon, Portugal.
- Causal Controversy at Le Mont St.-Michel. Robbert van Renesse. ACM Operating Systems Review 27, 2, (April 1993), 44-53.
- High Availability in a Real-Time System. Carlos Almeida, Bradford Glade, Keith Marzullo and Robbert van Renesse. ACM Operating Systems Review, 27, 2, (April 1993), 82-87.
- Fault-tolerant Programming using Process Groups. Robbert van Renesse and Kenneth Birman. IEEE Distributed Open Systems in Perspective. To appear in 1993.
- Fault-Tolerant Key Distribution (Preliminary Version). Michael Reiter, Kenneth P. Birman and Robbert van Renesse. January 1993. Revised version submitted for publication.

Describes the security architecture designed and built for the Horus system. Addresses several architectural issues about group programming vs. security.

- Priority Inversion and Its Prevention. Ozalp Babaoglu, Keith Marzullo and Fred B. Schneider. Journal of Real-time Systems 5, 285-303 (1993).

A formal definition of priority inversion and a set of protocols derived from this definition that completely avoid priority inversion.

- Integrating Runtime Consistency Models for Distributed Computing  
Kenneth Birman. To appear in Journal of Parallel and Distributed Computing.
- Monitoring and Controlling Distributed Applications using Lomita.  
Keith Marzullo and Ida Szafranska. Proceedings of the First Workshop  
on Systems Management, Los Angeles CA, May 1993.
- Obstacles to Software Reliability in Distributed Computing Environ-  
ments Kenneth Birman and Brad Glade. Submitted to IEEE Software,  
June 1993.